Workshop on Agentic AI Systems

Duration: 40 hours (05 days)

Table of Contents (TOC)

Module 1: Introduction to Agentic AI

- What is Agentic AI?
- Characteristics of Agentic AI Systems
- Use Cases and Applications
- Ethical Considerations and Risks

Module 2: Multi-Modal Models for Agentic AI

- Working with Language models
- Multi-modal models, gpt-4o
- OpenAI Whisper: Speech-to-text model (open-source), TTS model
- DALL-E Mini: Image generation and understanding (open-source alternative)
- Prompt Engineering

Module 3 LlamaIndex for RAG applications

- Introduction to indexing
- Llama index
- Understanding how indexing is done in Chroma
- Why Llama Index?
- Building blocks of Llama index
- Creating Indexes with Llama index
- Llama Parse

Module 4: RAG Pipelines Indexing PDFs, Docs, APIs into VectorDB and querying

- Introduction to Retrieval Augmented Generation (RAG)
- Using llms in llama index
- Document, PDFs, CSV, indexing
- Introduction to VectorDB
- Concept of Embedding, Retrieval, Chain and Agents using LangChain

- Lab: Build a Retrieval Augmented Generation (RAG) Application using LangChain

- Understanding and creating complex workflows in llamaindex

- Lab: Build an SQL database Agent

- Hands-On RAG Agent Build (Agentic RAG)

- Lab: Creating a chat bot which Converts Natural Language to SQL Query

Module 5 Building Agentic AI systems using LangChain

- Foundation of Function call

- Tool calling, Defining tools

- Creating custom tools

- LangChain for tool creation

- Configuring agent to return structured Response

Module 6: LangGraph for Creating Agentic AI Systems

- Understanding agents and creating agents in LangGraph

- LangChain: Building complex agent workflows

- Gradio: Creating interactive user interfaces

- Streamlit: Lightweight UIs for Agentic AI applications

- Multi agent systems using langgraph

Module 7: Agentic AI using Autogen

- Core idea of autogen

- Working with OpenAI models, ollama models, OS models

- Creating Assistant Agent with tools

- Creating multi-agent teams

- Implementing the reflection pattern

- Setting termination conditions

- Stopping, resuming, and aborting teams

- LangChain for Multi-Turn Conversations and Workflow Automation

- Integrating Gradio or Streamlit for Chatbot UIs

- Hands-On: Building a Task-Specific Chatbot

- Human-in-the-loop integration

Module 8: Evaluation and Optimization Using MLflow

- Introduction to MLflow for Model Tracking and Experimentation

- Logging Metrics, Parameters, and Artifacts for Agentic AI Models
- Model Versioning and Deployment with MLflow

Module 9: Training and Fine-Tuning Open-Source Models

- Customizing Models for Specific Agentic Tasks

- Tools for Model Fine-Tuning (**Hugging Face Trainer**, **PEFT**)

- Understanding about PEFT

- Understanding LORA

- Understandng QLORA

- Using LORA and QLORA to fine tune a model

Module 10: Ethical and Responsible Use of Agentic AI

- Bias Detection and Mitigation in Open-Source Models

- Transparency in Decision-Making Processes

- Safeguards Against Unintended Outcomes

Module 11: Hands-On Projects

- Building a Text-Based Autonomous Customer Support Agent

- Creating a Multi-Modal Autonomous Data Entry Assistant

- Developing a Inventory Management Agent

- Case Study: Building a Task Automation Pipeline

Module 12: Future Trends in Agentic AI

- Advances in Agent Architectures

- Evolution of Open-Source Model Ecosystems

- Cross-Domain Collaboration in Agentic AI